

# Yes, we CANN: Constrained Approximate Nearest Neighbors for local feature-based visual localization

**\*\*\* *Supplemental Material* \*\*\***

Dror Aiger      André Araujo      Simon Lynen

Google Research

{aigerd, andrearaujo, slynen}@google.com

## 1. Additional Qualitative Results

We include additional qualitative results in Figures 1,2,3,4 taken from all datasets, showing that CANN retrieves good results also in images with heavy occlusions. Cases like this, where there is only partial overlap between the query image and database images are very difficult for global features. We use HOW [2] for local features with both CANN-RG (ours) and ASMK [1]. The query image is on the left and the top 5 retrieved images are on the right. Our method retrieves all correct images, while other methods retrieve occasionally incorrect images ranked high among the top 5. We see that some global methods retrieve incorrect images due to scene clutter or high-frequency textures, while CANN provides diverse set of correct results. In several cases, we see that CANN+HOW outperforms ASMK+HOW. Retrieved images are marked red (bad) or green (good).

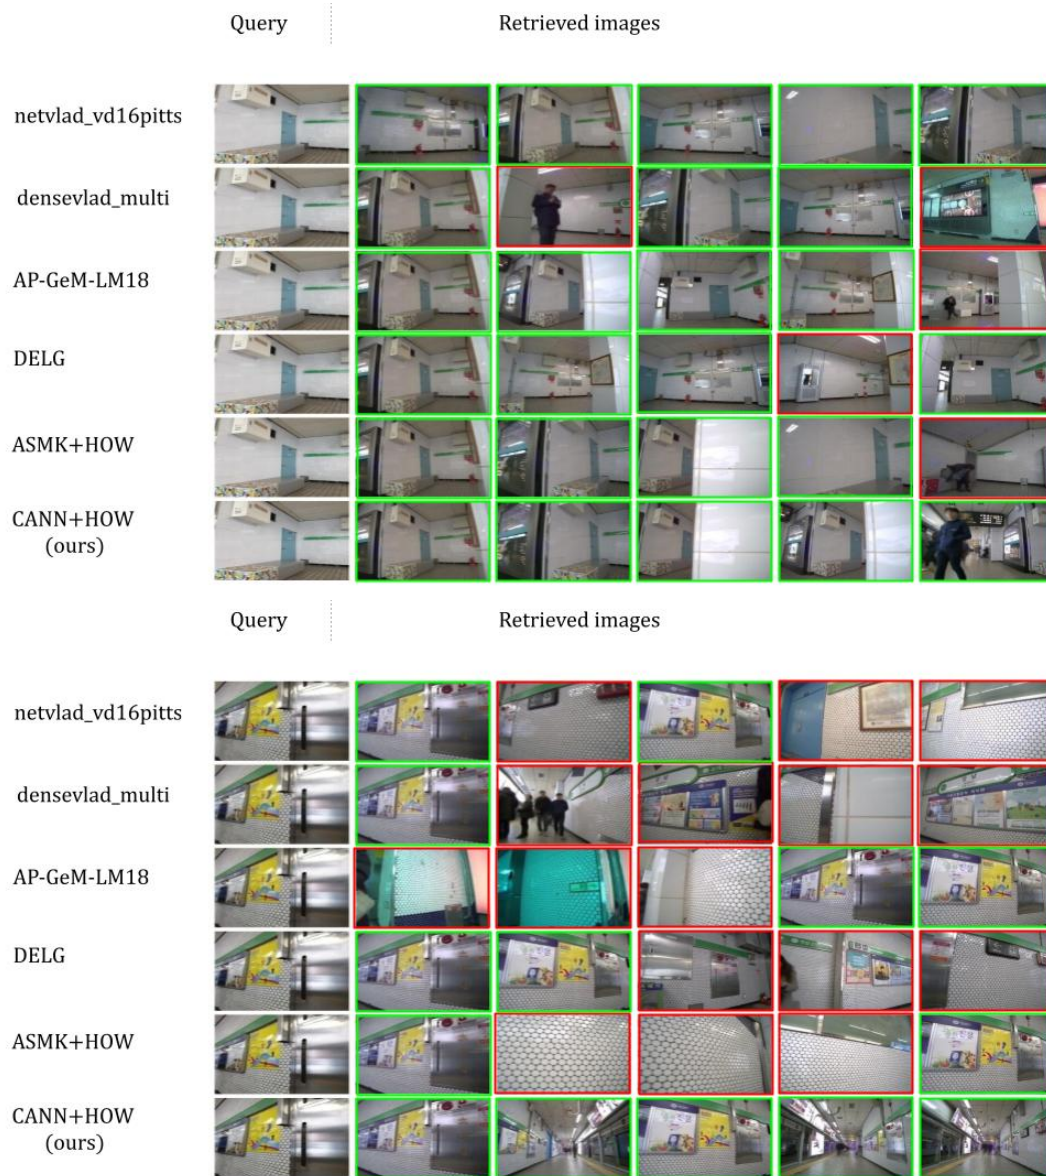


Figure 1: Baidu

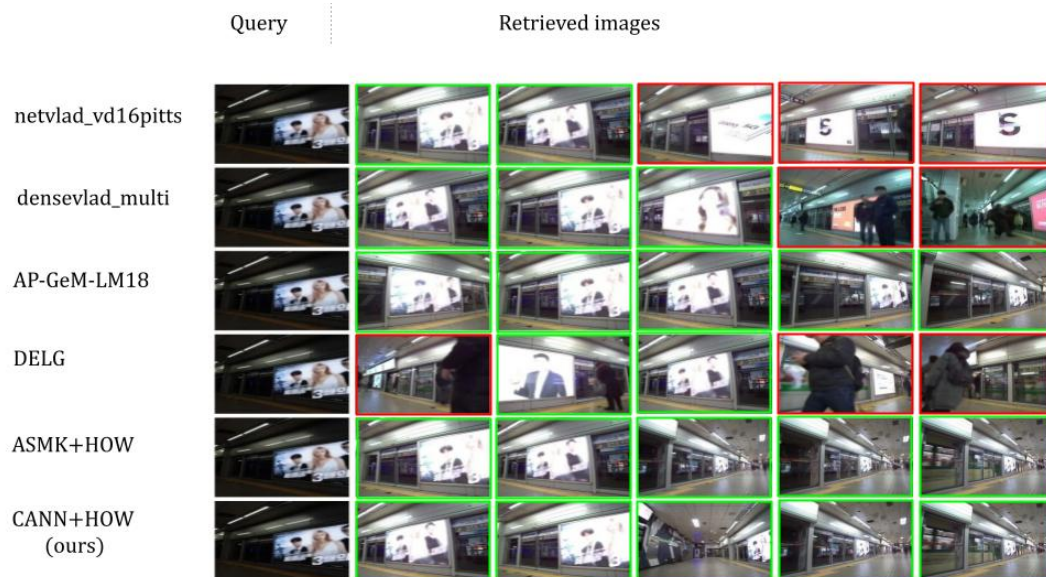


Figure 2: Baidu



Figure 3: Aachen

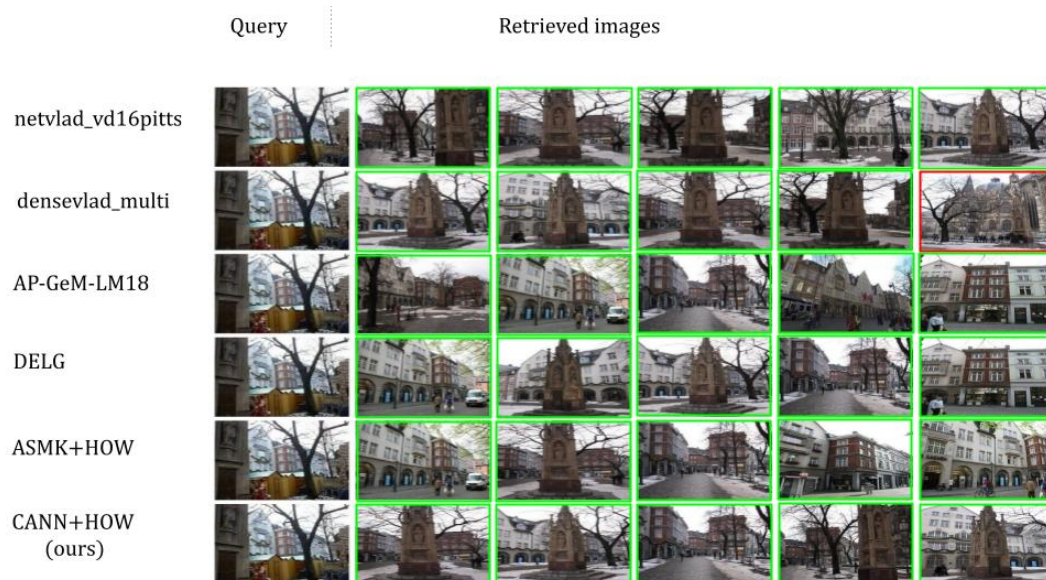


Figure 4: Aachen

## References

- [1] G. Tolias, Y. Avrithis, and H. Jegou. Image Search with Selective Match Kernels: Aggregation Across Single and Multiple Images. *IJCV*, 2015. 1
- [2] G. Tolias, T. Jenicek, and O. Chum. Learning and Aggregating Deep Local Descriptors for Instance-Level Recognition. In *Proc. ECCV*, 2020. 1